

PCT

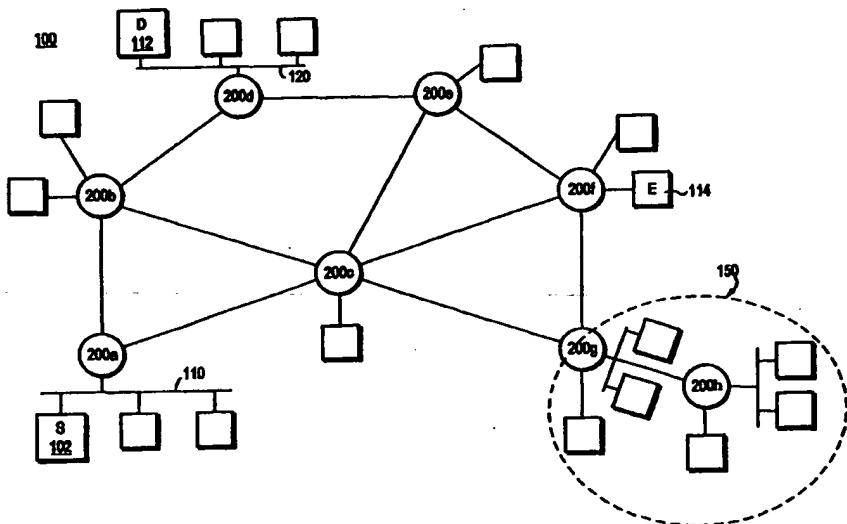
WORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau



INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : H04L 29/06		A1	(11) International Publication Number: WO 99/63727 (43) International Publication Date: 9 December 1999 (09.12.99)
(21) International Application Number:	PCT/US99/10272		(81) Designated States: AL, AU, BA, BB, BG, BR, CA, CN, CU, CZ, EE, GD, GE, HR, HU, ID, IL, IN, IS, JP, KP, KR, LC, LK, LR, LT, LV, MG, MK, MN, MX, NO, NZ, PL, RO, SG, SI, SK, SL, TR, TT, UA, UZ, VN, YU, ARIPO patent (GH, GM, KE, LS, MW, SD, SL, SZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).
(22) International Filing Date:	11 May 1999 (11.05.99)		
(30) Priority Data:	09/088,348	1 June 1998 (01.06.98)	US
(71) Applicant:	SUN MICROSYSTEMS, INC. [US/US]; 901 San Antonio Road, M/S PAL1-521, Palo Alto, CA 94303 (US).		
(72) Inventors:	GUPTA, Amit; 901 San Antonio Road, UMTV29-232, Palo Alto, CA 94303 (US). PERLMAN, Radia, Joy; 10 Huckleberry Lane, Acton, MA 01720 (US). CHIU, Dah-Ming; 736 Main Street, Acton, MA 01720 (US).		
(74) Agents:	BARBAS, Charles, J. et al.; Cesari and McKenna, LLP, 30 Rowes Wharf, Boston, MA 02110 (US).		
		Published	<i>With international search report. Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.</i>

(54) Title: DISTRIBUTED FILTERING AND MONITORING SYSTEM FOR A COMPUTER INTERNETWORK



(57) Abstract

A system efficiently distributes processing-intensive loads among a plurality of intermediate stations in a computer internetwork. The intermediate stations include routers, bridges, switches and/or firewalls configured with monitoring and filtering agents that communicate via a defined protocol to implement the system. Those stations configured with agents and having available resources cooperate to execute the loads which generally comprise verification operations on digital signatures appended to frame and/or packet traffic traversing paths of the computer internetwork. Techniques associated with the system are directed to efficiently detecting and filtering unauthorized traffic over portions of the internetwork protected as trust domains as well as unprotected portions of the internetwork.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav	TM	Turkmenistan
BF	Burkina Faso	GR	Greece		Republic of Macedonia	TR	Turkey
BG	Bulgaria	HU	Hungary	ML	Mali	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MN	Mongolia	UA	Ukraine
BR	Brazil	IL	Israel	MR	Mauritania	UG	Uganda
BY	Belarus	IS	Iceland	MW	Malawi	US	United States of America
CA	Canada	IT	Italy	MX	Mexico	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NE	Niger	VN	Viet Nam
CG	Congo	KR	Kenya	NL	Netherlands	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NO	Norway	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's	NZ	New Zealand		
CM	Cameroon		Republic of Korea	PL	Poland		
CN	China	KR	Republic of Korea	PT	Portugal		
CU	Cuba	KZ	Kazakhstan	RO	Romania		
CZ	Czech Republic	LC	Saint Lucia	RU	Russian Federation		
DE	Germany	LI	Liechtenstein	SD	Sudan		
DK	Denmark	LK	Sri Lanka	SE	Sweden		
EE	Estonia	LR	Liberia	SG	Singapore		

DISTRIBUTED FILTERING AND MONITORING SYSTEM FOR A COMPUTER INTERNETWORK

FIELD OF THE INVENTION

The present invention relates generally to computer networks and, more specifically, to a technique for distributing processing loads among intermediate stations of a computer internetwork.

BACKGROUND OF THE INVENTION

Communication in a computer internetwork involves the exchange of data between two or more entities interconnected by communication media. The entities are typically software programs executing on hardware computer platforms, such as end stations and intermediate stations. In particular, communication software executing on the stations correlate and manage data communication with other stations. The stations typically communicate by exchanging discrete packets or frames of data according to predefined protocols. A protocol, in this context, consists of a set of rules defining how the stations interact with each other.

Examples of intermediate stations include a router, a switch and a bridge, each of which generally comprises a processor, such as a central processing unit (CPU), and port interfaces for interconnecting a wide range of communication links and subnetworks. Broadly stated, the router implements network services such as route processing, path determination and path switching functions. The route processing function determines optimal routing for a packet based, e.g., on a destination address stored in a packet header, such as an Internet protocol (IP) header, whereas the path switching function allows a router to accept a packet on one interface and forward it on a second interface. The path determination, or forwarding decision, function selects the most appropriate interface for forwarding a packet. The switch, on the other hand, provides the basic functions of a bridge including filtering of data traffic by medium access control (MAC) address and "learning" of a MAC address based upon a source MAC address of a frame; in addition, the switch may implement forwarding decision and path switching operations of a router.

-2-

Often it may be desirable for the intermediate station, hereinafter referred to as a switch, to perform additional processing-intensive operations on a stream of data traffic. An example of such processing may involve policing of bandwidth utilization policies, such as usage parameter control or network parameter control policies. Violations of these policies typically occur when a transmitting station ("sender") attempts to utilize more bandwidth than it is entitled to under the terms of the policies. Hence, it may be desired to have the switch record, e.g., all frame traffic of a particular type and within a particular time span that are transmitted from a particular sender to ensure that the sender does not exceed a particular threshold (e.g., 5 frames a second). Enforcement of the policies results in the switch discarding (filtering) those frames in violation of the bandwidth constraints. However, the accounting, auditing and/or record keeping operations needed to ensure that the bandwidth constraints are observed may tax the CPU resources of the switch.

Another example of desired processing-intensive operations for a switch involves checking for invalid digital signatures on packets received at the switch. Digital signatures are generally used in accordance with a well-known cryptographic technique for performing remote authentication known as public key cryptography. In this method of secure communication, each entity has a public encryption key and a private encryption key, and two entities can communicate knowing only each other's public keys and, of course, their own private keys. To prove to a recipient of information that the sender is who it purports to be, the sender digitally encodes ("signs") the information with its private key. If the recipient can decode ("verify") the information, it knows that the sender has correctly identified itself.

It may be desirable for the switch to check digital signatures on packets to verify that the packets originate from a sender that is authorized to use the bandwidth of the network. Here, the authorized senders have access to a group private key and the switch has access to, or acquires, the complementary group public key. The authorized senders thus sign their packets with the group private key and the switch verifies those signatures upon receiving the packets. Those packets having unauthorized signatures may be filtered by the switch.

However, it typically takes longer for a switch to verify a signature than it takes to receive a small packet at wire speed, so an unauthorized sender could mount a "denial-of-service" attack by transmitting a stream of invalidly signed packets. Such an attack becomes particularly burdensome when the switch attempts to verify digital signatures on 5 multicast packets. Checking of each packet at the switch consumes valuable CPU resources which typically cannot be spared because of their occupation with conventional switching operations. Each switch is generally independently responsible for verifying packets of a traffic stream and, as such, a single switch may be unable to keep up with the traffic of such an attack. Therefore, the present invention is directed to alleviating the processing-intensive 10 burden placed on a switch of a computer internetwork.

SUMMARY OF THE INVENTION

The invention relates to a load distribution system for efficiently distributing processing-intensive loads among a plurality of intermediate stations in a computer internetwork. The intermediate stations include routers, bridges and/or switches configured with 15 monitoring and filtering agents that communicate via a defined protocol to implement the system. According to the invention, those stations configured with agents and having available resources cooperate to execute the loads which generally comprise verification operations on digital signatures appended to frame and/or packet traffic traversing paths of the computer internetwork. The novel system includes various techniques that are directed to 20 efficiently detecting and filtering unauthorized traffic, hereinafter referred to as packets, over portions of the internetwork protected as trust domains as well as unprotected portions of the internetwork, such as the Internet.

In a first technique of the inventive system, multiple switches share the verification load by independently processing a random selection of packets. Most senders of packets 25 generally do not transmit unauthorized packets over the internetwork; therefore, rather than process (check) each packet, the first technique employs a plurality of switches to "spot-check" a fraction of the packets. Notably, there is no need to coordinate processing among

the switches. A second technique, however, contemplates examination of each packet, but in a manner that avoids redundancy. A hash function is invoked to assign packet checking responsibility to each switch configured with the monitoring and filtering agent. Communication among the agents resident on switches of a path traversed by the packets enables distinguishing of responsibility based on, e.g., even/odd destination addresses of the packets.

If apportioning of the processing load according to the second technique results in unbalanced division of work among the switches, a sophisticated hash function is employed that divides the packets into substantially more "buckets" than the previous hash function.

10 According to this third inventive technique, each switch is assigned an equal number of buckets and if any switch is burdened with excessive processing responsibility, additional buckets are exchanged among the switches.

In a trust domain, an enterprise *trusts* only those switches within its domain, yet packets are often exchanged with an "outside" network, such as the Internet. A fourth technique of the inventive load distribution system specifies the use of a flag in the header of the packet to indicate whether the packet has been verified by a trusted switch. The flag may be contained with an unused field of the header or, alternatively, it may be part of a mini-header added to the packet for use within the trust domain.

A fifth technique is directed to changing the fraction of packets that are "spot-checked" on a per flow basis. A flow is a function of information stored in a packet header such as, in the case of an Internet protocol (IP) packet, source and destination IP address, source and destination transmission control protocol (TCP) ports and flow label. This aspect of the invention may be advantageously employed with the first technique such that once an unauthorized packet is identified via random spot-checking, a more careful examination of the source stream is conducted. Various approaches to changing the fraction of packets that are spot-checked for each flow are described further herein.

A sixth technique of the inventive system is directed to distributed filtering of packets using non-optimal routing. Since there may be only certain switches in the internetwork

configured with the agents, this technique specifies changing an optimal path taken by a packet so it traverses switches having packet-verification capability.

Advantageously, the invention described herein distributes data traffic processing loads among intermediate stations of a computer internetwork in an efficient manner. More specifically, the invention provides a distributed filtering mechanism configured to avoid "denial-of-service" attacks on a switch of a computer internetwork.

BRIEF DESCRIPTION OF THE DRAWINGS

The above and further advantages of the invention may be better understood by referring to the following description in conjunction with the accompanying drawings in which like reference numbers indicate identical or functionally similar elements:

Fig. 1 is a schematic block diagram of a computer internetwork comprising a collection of interconnected communication media attached to a plurality of stations, including intermediate stations;

Fig. 2 is a schematic block diagram of an intermediate station, such as a network switch, configured to operate in accordance with a load distribution system for efficiently distributing processing-intensive loads among similarly-configured switches of the computer internetwork of Fig. 1;

Figs. 3 is a schematic diagram illustrating a hash function technique for use in accordance with the inventive load distribution system;

Fig. 4 is a schematic diagram illustrating a more sophisticated hash function technique for use in accordance with the inventive load distribution system;

Figs. 5A and 5B are schematic diagrams illustrating a trusted flag technique for use in accordance with the inventive load distribution system; and

Fig. 6 is a schematic diagram illustrating a flow altering technique for use in accordance with the inventive load distribution system.

DETAILED DESCRIPTION OF AN ILLUSTRATIVE EMBODIMENT

Fig. 1 is a schematic block diagram of a computer internetwork 100 comprising a collection of interconnected communication media attached to a plurality of stations. The communication media may be configured as local area networks (LANs) 110 and 120, although other media configurations such as point-to-point network links may be advantageously employed. The stations are typically computers comprising source (S) and destination (D) end stations 102, 112, such as personal computers or workstations, and intermediate stations 200a-h such as routers, bridges, switches and/or firewalls. Communication among the stations is typically effected by exchanging discrete data frames or packets between the communicating nodes according to a predefined protocol, such as the Internet protocol (IP), Internet Packet Exchange protocol, AppleTalk protocol or DECNet protocol.

Fig. 2 is a schematic block diagram of an intermediate station, hereinafter referred to as a switch 200, configured to operate in accordance with the invention, although it should be noted that other intermediate stations, such as a router or bridge, may be advantageously used with the present invention. The switch generally comprises a plurality of interconnected elements, such as a processor 202, a memory 204 and a network adapter 206 having a plurality of port interfaces 208a-c. The memory may comprise storage locations addressable by the processor and adapter for storing software programs and data structures associated with the inventive load distribution techniques. The processor is typically a central processing unit (CPU), although it may generally comprise processing elements or logic configured to execute the software programs and manipulate the data structures. It will be apparent to those skilled in the art that other processor and memory means, including various computer readable media, may be used for storing and executing program instructions pertaining to the techniques described herein.

An operating system 210, portions of which are typically resident in memory and executed by the CPU, functionally organizes the station by, *inter alia*, invoking network operations in support of software processes executing on the station. An example of such a

software process is a conventional monitoring and filtering agent 220 that configures the switch for communication with agents of similarly-configured switches to thereby enable implementation of the inventive load distribution system described herein. Preferably, the filtering agents communicate via a conventional network management protocol, such as a

5 Simple Network Management Protocol (SNMP).

The inventive load distribution system efficiently distributes processing-intensive loads among a plurality of switches configured with the monitoring and filtering agents. Specifically, those switches configured with the agents and having available CPU resources cooperate to execute the loads which generally comprise verification operations on digital 10 signatures appended to frame and/or packet traffic traversing paths of the computer internetwork. To that end, the novel system includes various techniques that are directed to efficiently detecting and filtering unauthorized traffic, hereinafter referred to as packets, over portions of the internetwork protected as trust domains as well as unprotected portions of the internetwork, such as the Internet.

15 According to a first technique, multiple switches share the verification load by independently processing a random selection of packets. Most senders of packets generally do not transmit unauthorized packets over the internetwork; therefore, rather than process (check) each packet, the first technique employs a plurality of switches to "spot-check" a fraction of the packets. Notably, there is no need to coordinate processing among the 20 switches.

For example refer to Fig. 1 wherein two switches 200a,b of a computed path are dynamically assigned monitoring and verification responsibilities. If each of the switches randomly check $1/n$ of the packet traffic, then only $1/n^{**2}$ packets are checked twice. If n is reasonably large (e.g., 5 or more) then collectively the switches check roughly $2/n$ of the 25 packets. This technique provides a probabilistic approach where there is no requirement to process each packet, but where there is some value associated with checking occasional packets.

According to a second technique of the system, a hash function is invoked by the switch to assign packet-checking responsibility among switches configured to operate in accordance with the invention. Here, each packet is examined by the switches and assigned to a responsible switch based on a *hash* of the packet in a manner that avoids redundancy.

5 A hash is a function that generally reduces a large quantity to smaller quantities of "buckets" that requires considerably less CPU resources to compute than digital signature verification. In general, any field/aspect of a packet can be hashed if it facilitates distinguishing the packets.

For example, agents 220 resident on switches 200a,b of a path traversed by the
10 packets communicate to distinguish responsibility based on addresses of the packets. Fig. 3
is a schematic diagram illustrating a reasonable hash function 300 based on a destination
address 310 of a header 325 of packet 320. According to the technique, all even destination
addresses are assigned to bucket 330 and all odd destination addresses are assigned to
15 bucket 340. Thereafter, each bucket is assigned to a respective switch 200a,b. Another
example of a reasonable hash function relates to distinguishing data packets by even/odd
sequence numbers.

If apportioning of the processing load according to the second technique results in unbalanced division of work among the switches because, e.g., there are substantially more even addresses than there are odd addresses, a more sophisticated hash function may be
20 employed that divides the packets into substantially more buckets than the previous hash function. According to this third inventive technique, each switch is assigned an equal number of buckets and if any switch is burdened with excessive processing responsibility, additional buckets are exchanged among the switches.

Fig. 4 is a schematic diagram of an example of such a hash function 400 involving
25 use of, e.g., four (4) least significant bits (LSB 415) of a source address 410 of a packet
420. These LSBs translate into sixteen (16) different values 430a-p, each of which is assigned to a bucket 440a-p. The packets are then apportioned among the switches 200a,b based on, e.g., even/odd bucket values. If it turns out that more packets are processed by

the switch assigned the even values or if the switch assigned the odd values has additional CPU capacity, then the latter switch can be assigned additional packets, such as ones with the values 4 and 12. It is apparent that more than 16 buckets may be realized by employing additional LSBs and, of course, the same technique could apply to most significant bits of the packets or, for that matter, anything that divides the packets into buckets. This technique thus enables even distribution of the load or, alternatively, intentional uneven distributions thereof, if necessary, among the switches.

Essentially, this third hash function technique represents a multi-level approach wherein the load is initially divided by a predetermined criterion obtained as a result of decoding predetermined address bits (first level) and then, if it is apparent over a period of time that this initial apportionment is ineffective, a second level is invoked to balance the load based on such empirical data. For example as traffic characteristics of packet transmission over the network change, distribution of the load may also change.

Often a portion of the internetwork is protected as a *trust domain*, an example of which may be a corporate intranet of a particular enterprise which is schematically shown as domain 150 in Fig. 1. Typically, the enterprise *trusts* only those switches within its domain, yet packets are usually exchanged between the trust domain 150 and an "outside" network, such as the Internet, through a switch 200g configured as a firewall. Figs. 5A and 5B are schematic diagrams illustrating a fourth technique of the inventive load distribution system wherein a flag (F) 515 in a header of a packet 500 is used to indicate whether the packet has been verified by a trusted switch 200g,h. According to the inventive technique, the flag 515 may be contained with an unused field 510 of the header or, alternatively, it may be part of a mini-header 530 added to the packet 500 for use within the trust domain.

For most protocols, there are spare bits in the header that are not generally used, such as the bits contained in a "type of service" field on an IPv6 header. Within a trust domain 150, one of these unused bits can be designated as the flag for use by the trusted switches 200g,h according to an algorithm of the present invention. That is, the flag 515 is defined as initialized to zero so that when a packet 500 is received at the firewall 200g from

-10-

the Internet, the firewall ensures that the flag is cleared; this indicates that the packet was received from outside the trust domain and that it has yet to be checked by a trusted switch. Even if the packet was checked prior to traversing the firewall, the flag 515 is cleared because the enterprise does not trust any switch outside of the trust domain to perform the verification operations.

If there are no spare bits in the packet header 525, the firewall may alternately add a mini-header 530 to the packet for use within the trust domain to carry the flag. Mini-headers are often appended to packets, such as in the case of virtual local area networks (VLANs) and tag switching. The appended mini-header has a length of at least one bit for indicating packet verification (or lack thereof) by a trusted switch.

Thus, according to the fourth technique, each incoming packet received at the trust domain 150 from outside of that domain is marked indicating whether the packet needs to be checked by a trusted switch. For example, setting of the flag denotes that no other switch need examine the packet. Any switch within the trust domain configured to check packets and having available CPU capacity may examine the packet and thereafter set the flag indicating that it has been verified. This aspect of the inventive technique allows "upstream" switches to check unverified packets depending upon their available capacity (and time), while allowing "downstream" switches having such available capacities to process any remaining unverified packets.

Moreover, use of the "trusted flag" technique may obviate the need for the second and third techniques described herein because a trusted switch configured with the present invention will perform a packet verification procedure only if it has available resources. The fourth technique may be used, however, with the first random checking technique where the need to verify a packet is evident from the state of the defined flag 515. In other words, marking the packets according to the fourth technique is useful in connection with the first technique to eliminate duplicate packet checking, but it is not generally useful with the second or third techniques because the switches have been assigned responsibility for certain packets.

-11-

Fig. 6 is a schematic diagram illustrating a fifth technique of the invention wherein a fraction of packets that are spot-checked may be altered on a per flow basis. A flow is a function of information stored in a header of a packet; in the case of a transmission control protocol/Internet protocol (TCP/IP) packet 600, the flow may comprise source and destination addresses 612, 614 of an IP header 610, source and destination ports 622, 624 of a TCP header 620 and a flow label 615, which is a 28-bit field of an IPv6 packet header. This "flow" technique may be advantageously employed with the first technique described herein such that once an unauthorized packet is identified via random spot-checking, a more careful examination of the source stream is conducted. Various approaches to changing the fraction of packets that are spot-checked for each flow are described further herein.

Specifically, a fraction (e.g., one out of every thousand) of the packets received from a particular source (identified by an address prefix of a particular address block) is initially checked and as long as no unauthorized packets are identified, the fractional spot-checking ratio is maintained. Yet once an unauthorized packet is identified from that source, a more careful examination of the stream is conducted. For example, if an unauthorized packet is discovered having a source prefix 685*, it is assumed that a sender at 685* is attempting to use bandwidth on the network in violation of a prescribed policy. Accordingly, this aspect of the load distribution system provides a number of approaches to changing the fraction of packets that are spot-checked for each flow.

(a) The switch can be manually configured to check a certain (probably increased) fraction of packets based on a particular flow. For example, for packets having particular destination port 624 or a particular prefix of a source address 612, check fifty percent of the packets rather than one percent.

(b) Initially only a small fraction of packets (e.g., one out of every thousand) is checked but once a packet with a particular flow is identified as having an unauthorized signature on it, a larger fraction of packets with that flow is thereafter checked by the switch. Note that for this approach, the switch that notices the problem changes the fraction itself.

-12-

(c) This approach is generally the same as (b) except that the switch is informed of the unauthorized packet by a trusted switch or a system administrator.

(d) Although the switch does not find any unauthorized packets, it does detect a change in the traffic pattern, e.g., a significant increase in packet transfer rate from a particular source. This may be an indication that a sender is attempting to attack the network by inserting a plurality of unauthorized packets into a stream. The switch thus increases the fraction of packets that are inspected based on a particular flow.

A sixth technique of the invention is directed to distributed filtering of packets using non-optimal routing. If there are only certain switches configured with agents that are capable of checking packets, this technique specifies changing the optimal path calculated for a packet so that it will traverse switches that have packet-verification capabilities. For example and referring to Fig. 1, assume a minimal path to destination 112 from end station (E) 114 encompasses switches 200f,e,d; however none of these switches are capable of examining a packet to verify an appended digital signature. Therefore, according to this inventive technique, the packet is forwarded along a non-optimal path that includes switches 200b,c, each of which is configured to check the packet.

In the case of a connection-oriented network, such as an ATM network where virtual circuits are established, there may be two different virtual circuits to a destination: one representing the shortest path but which encompasses no additional switches capable of packet-verification processing and another that, although non-optimal, includes switches that can process the packet. If a particular switch has the resources (time or capacity) to process the packet, it checks the packet and if verified, sends it over the minimal path. If the switch does not have the resources to process the packet, it sends the packet over the non-optimal path in hopes that a downstream switch will check and verify it.

For a connectionless network, each switch sends the packet to a neighbor. In the case of an IP packet, each switch forwards the packet to a next hop and the next hop makes an independent decision as to the direction of the packet. That is, the next hop could decide

-13-

that another switch having the capability of checking for unauthorized packets should be included in the path, even though that path is not the optimal route to the destination.

While there has been shown and described an illustrative embodiment for efficiently distributing processing-intensive loads among a plurality of intermediate stations in a computer internetwork, it is to be understood that various other adaptations and modifications may be made within the spirit and scope of the invention. For example, in an alternative embodiment of the fourth inventive technique, unverified packet traffic may be treated as lower priority than verified packet traffic. Recall that a reason for checking the digital signatures of the packets is to ensure that unauthorized packets do not consume bandwidth in the network. When an unauthorized packet is identified, it is discarded. There may be a situation where the trusted switches do not have the resources (capacity and/or time) to check all of the packets. Those packets that have been checked are, according to this alternate embodiment, distinguished with an asserted predetermined flag 515 and are treated as high priority packets.

Specifically, if firewall 200g at an entry point of the trust domain 150 cannot keep up with the incoming traffic, the interior trusted switches 200h need not verify any unverified packets, but rather may choose to treat those packets as low priority. That is, those packets that are verified are placed on high priority queues (H 218 of Fig. 2) within a switch 200 and the unverified packets are placed on low priority queues (L 214). The low priority packets are then prone to "dropping" (discarding) if the trusted region exceeds a certain bandwidth utilization. Thus even though there is not enough CPU capacity to check all the packets, there is still some benefit to checking a fraction of the packets.

The foregoing description has been directed to specific embodiments of this invention. It will be apparent, however, that other variations and modifications may be made to the described embodiments, with the attainment of some or all of their advantages. Therefore, it is the object of the appended claims to cover all such variations and modifications as come within the true spirit and scope of the invention.

What is claimed is:

CLAIMS

- 1 1. A method for efficiently distributing processing-intensive loads among a plurality of intermediate stations in a computer internetwork, the method comprising the steps of:
 - 3 configuring at least one intermediate station with a monitoring and filtering agent
 - 4 process adapted to execute the loads on data traffic traversing paths of the computer internetwork; and
 - 6 at the configured intermediate station, independently processing a selection of data traffic to thereby share the loads among the intermediate stations.
- 1 2. The method of Claim 1 wherein the data traffic comprises packets and wherein processing of the loads comprises verification operations on digital signatures appended to the packets.
- 1 3. The method of Claim 1 wherein the data traffic comprises packets and wherein the selection of packets processed by each intermediate station is randomly assigned according to a fractional spot-checking function.
- 1 4. The method of Claim 1 wherein the data traffic comprises packets and wherein the selection of packets processed by the configured intermediate station is assigned according to a hash function.
- 1 5. The method of Claim 4 further comprising the step of invoking the hash function to distinguish the packets based on contents of a field of each packet.
- 1 6. The method of Claim 5 further comprising the steps of, wherein the contents comprise a destination address:
 - 3 apportioning the packets according to even and odd destination addresses; and

-15-

4 assigning packets having even destination addresses to one intermediate station and
5 packets having odd destination addresses to another intermediate station.

1 7. The method of Claim 5 further comprising the steps of, wherein the contents comprise a
2 sequence number:

3 apportioning the packets according to even and odd sequence numbers; and
4 assigning packets having even sequence numbers to one intermediate station and
5 packets having odd sequence numbers to another intermediate station.

1 8. The method of Claim 5 further comprising the steps of, wherein the contents comprise an
2 address:

3 decoding predetermined address bits into distinct values;
4 apportioning the packets according to the distinct values;
5 assigning packets to the intermediate stations based on the distinct values; and
6 reassigning certain of the previously-assigned packets to certain of the intermediate
7 stations to thereby balance the load among the stations.

1 9. The method of Claim 3 further comprising the steps of, wherein the data traffic comprises
2 packets and wherein processing of the loads comprises verification operations on digital
3 signatures appended to the packets:

4 in response to identifying an unauthorized packet, filtering the unauthorized packet;
5 and
6 altering the random fractional spot-checking of packets by each intermediate station
7 on a per flow basis, wherein a flow comprises source and destination addresses, source and
8 destination ports and a flow label of a packet.

1 10. The method of Claim 9 wherein the step of altering comprises the step of spot-checking
2 an increased fraction of the packets.

1 11. The method of Claim 9 wherein the step of altering comprises the step of manually con-
2 figuring the intermediate station to spot-check an increased fraction of the packets based on
3 the flow.

1 12. The method of Claim 3 further comprising the step of, wherein the data traffic com-
2 prises packets:

3 in response to identifying a change in traffic pattern, altering the random fractional
4 spot-checking of packets by each intermediate station on a per flow basis, wherein a flow
5 comprises source and destination addresses, source and destination ports and a flow label of
6 a packet.

1 13. The method of Claim 2 further comprising the step of, wherein a portion of the inter-
2 network is protected as a trust domain having trusted switches:

3 employing a flag within a header of a packet to indicate whether the packet has been
4 verified by a trusted switch configured with the monitoring and filtering agent.

1 14. The method of Claim 13 wherein the flag may be contained within one of an unused
2 field of the header and a mini-header appended to the packet.

1 15. The method of Claim 14 further comprising the step of, if a state of the flag indicates
2 that the packet has not been verified by an upstream trusted switch along a path of the com-
3 puter internetwork:

4 enabling a downstream trusted switch to process the unverified packet depending
5 upon its available capacity.

1 16. The method of Claim 2 further comprising the steps of:
2 calculating an optimal path for the packets over the computer internetwork;

3 determining whether intermediate stations located along the optimal path are con-
4 figured with the monitoring and filtering agents; and
5 changing the optimal path of the packets to a non-optimal path that includes addi-
6 tional intermediate stations configured with the monitoring and filtering agents adapted to
7 execute the verification operations on the packets.

1 17. The method of Claim 1 wherein the intermediate station comprises a switch.

1 18. The method of Claim 1 wherein the intermediate station comprises a router.

1 19. The method of Claim 1 wherein the intermediate station comprises a bridge.

1 20. The method of Claim 1 wherein the intermediate station comprises a firewall.

1 21. The method of Claim 3 wherein the intermediate station comprises a firewall.

1 22. Apparatus for efficiently distributing processing-intensive loads among a plurality of in-
2 termediate stations in a computer internetwork, the apparatus comprising:

3 means for configuring at least one intermediate station with a monitoring and filter-
4 ing agent process adapted to execute the loads on data traffic traversing paths of the com-
5 puter internetwork; and

6 means for independently processing a selection of data traffic at the configured in-
7 termediate station to thereby share the loads among the intermediate stations.

1 23. A computer readable medium containing executable program instructions for efficiently
2 distributing processing-intensive loads directed to verification operations on digital signa-
3 tures appended to packets transferred among a plurality of intermediate stations in a com-

-18-

4 computer internetwork, the executable program instructions comprising program instructions
5 for:

6 configuring at least one intermediate station with a monitoring and filtering agent
7 process adapted to execute the loads on the packets traversing paths of the computer inter-
8 network; and

9 at the configured intermediate station, independently processing a selection of the
10 packets assigned to the station according to a hash function that enables checking of the
11 digital signatures to identify one of authorized and unauthorized packets, thereby enabling
12 sharing of the loads among the intermediate stations.

1 24. The computer readable medium of Claim 23 wherein said executable program instruc-
2 tions further comprise program instructions for invoking the hash function to distinguish the
3 packets based on contents of a field of each packet.

1 25. The computer readable medium of Claim 24 wherein said executable program instruc-
2 tions further comprise program instructions for, wherein the contents comprise an address:

3 decoding predetermined address bits into distinct values;
4 apportioning the packets according to the distinct values;
5 assigning packets to the intermediate stations based on the distinct values; and
6 reassigned certain of the previously-assigned packets to certain of the intermediate
7 stations to thereby balance the load among the stations.

1 26. A computer data signal embodied in a carrier wave and representing sequences of in-
2 structions for efficiently distributing processing-intensive loads directed to verification op-
3 erations on digital signatures appended to packets transferred among a plurality of interme-
4 diate stations in a computer internetwork, the instructions comprising instructions for:

5 configuring at least one intermediate station with a monitoring and filtering agent
6 process adapted to execute the loads on the packets traversing paths of the computer inter-
7 network; and

8 at the configured intermediate station, independently processing a selection of the
9 packets assigned to the station according to a hash function that enables checking of the
10 digital signatures to identify one of authorized and unauthorized packets, thereby enabling
11 sharing of the loads among the intermediate stations.

1 27. The computer data signal of Claim 26 wherein the selection of packets processed by
2 each intermediate station is randomly assigned according to a fractional spot-checking
3 function to enable checking of the digital signatures to identify one of authorized and un-
4 authorized packets.

1 28. The computer data signal of Claim 27 wherein the instructions further comprise instruc-
2 tions for:

3 in response to identifying an unauthorized packet, filtering the unauthorized packet;
4 and
5 altering the random fractional spot-checking of packets by each intermediate station
6 on a per flow basis.

1 29. The computer data signal of Claim 28 wherein the instructions for altering further com-
2 prise instructions for spot-checking an increased fraction of the packets.

1 30. A system for efficiently distributing processing-intensive loads among a plurality of in-
2 termediate stations in a computer internetwork, the system comprising:

3 a plurality of memory devices containing software programs organized as monitor-
4 ing and filtering agents adapted to execute the loads on data traffic traversing paths of the
5 computer internetwork; and

6 a plurality of processing elements coupled to respective ones of the memory devices,
7 each processing element configured to execute a respective agent to independently process
8 a selection of data traffic to thereby share the loads among the intermediate stations.

-20-

- 1 31. The system of Claim 30 wherein the data traffic comprises packets and wherein processing of the loads comprises verification operations on digital signatures appended to the packets.
- 1 32. The system of Claim 31 further comprising, wherein a portion of the internetwork is protected as a trust domain having trusted switches:
 - 3 a flag structure contained within a header of a packet to indicate whether the packet
 - 4 has been verified by a trusted switch configured with the monitoring and filtering agent.
- 1 33. The system of Claim 32 wherein the flag structure may be contained within one of an unused field of the header and a mini-header appended to the packet.
- 1 34. The system of Claim 33 wherein the intermediate station comprises one of a switch and a router.

1/6

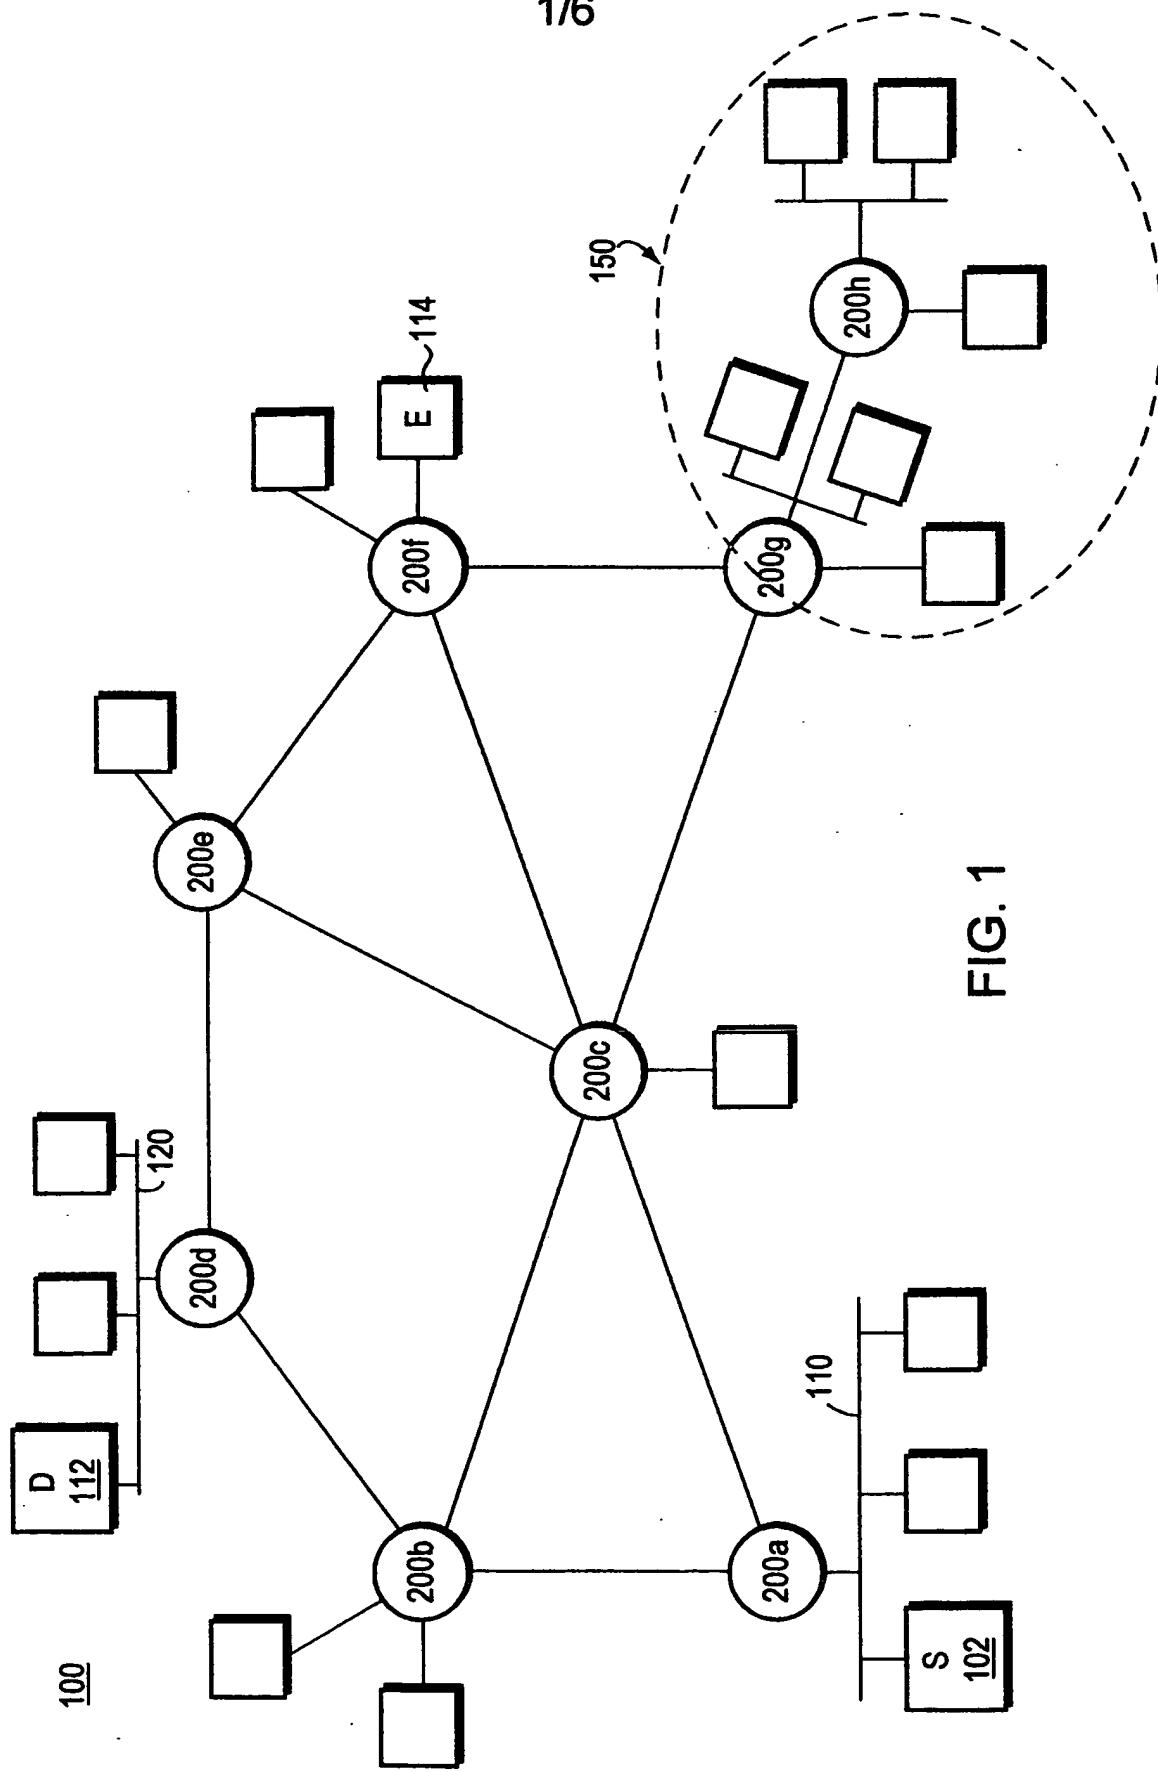


FIG. 1

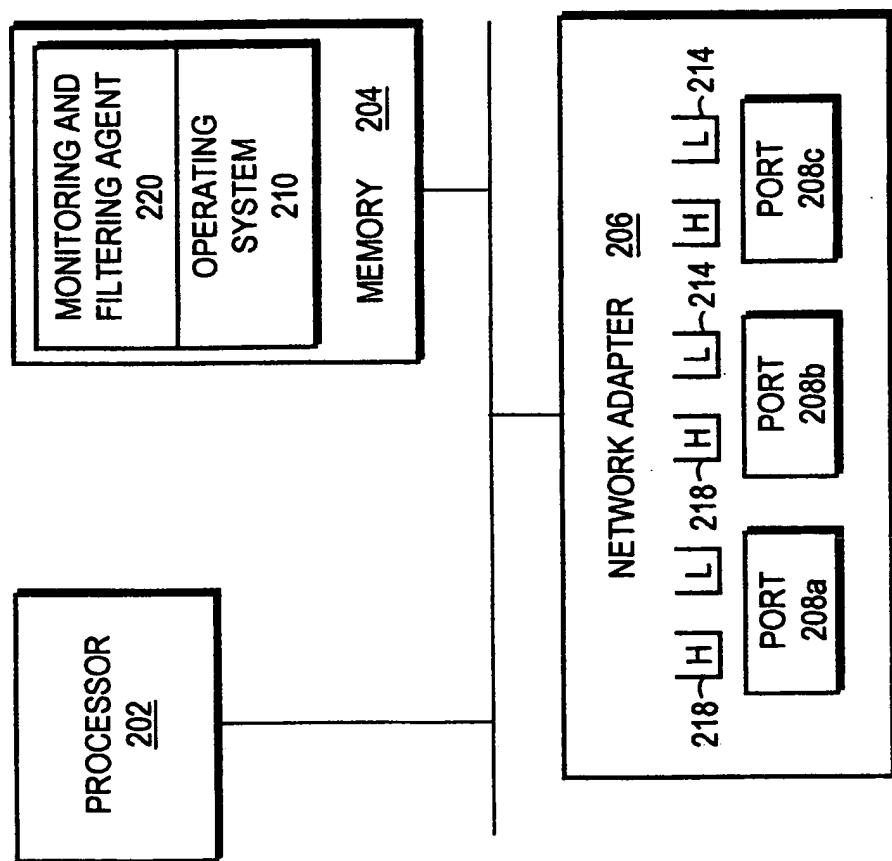
200

FIG. 2

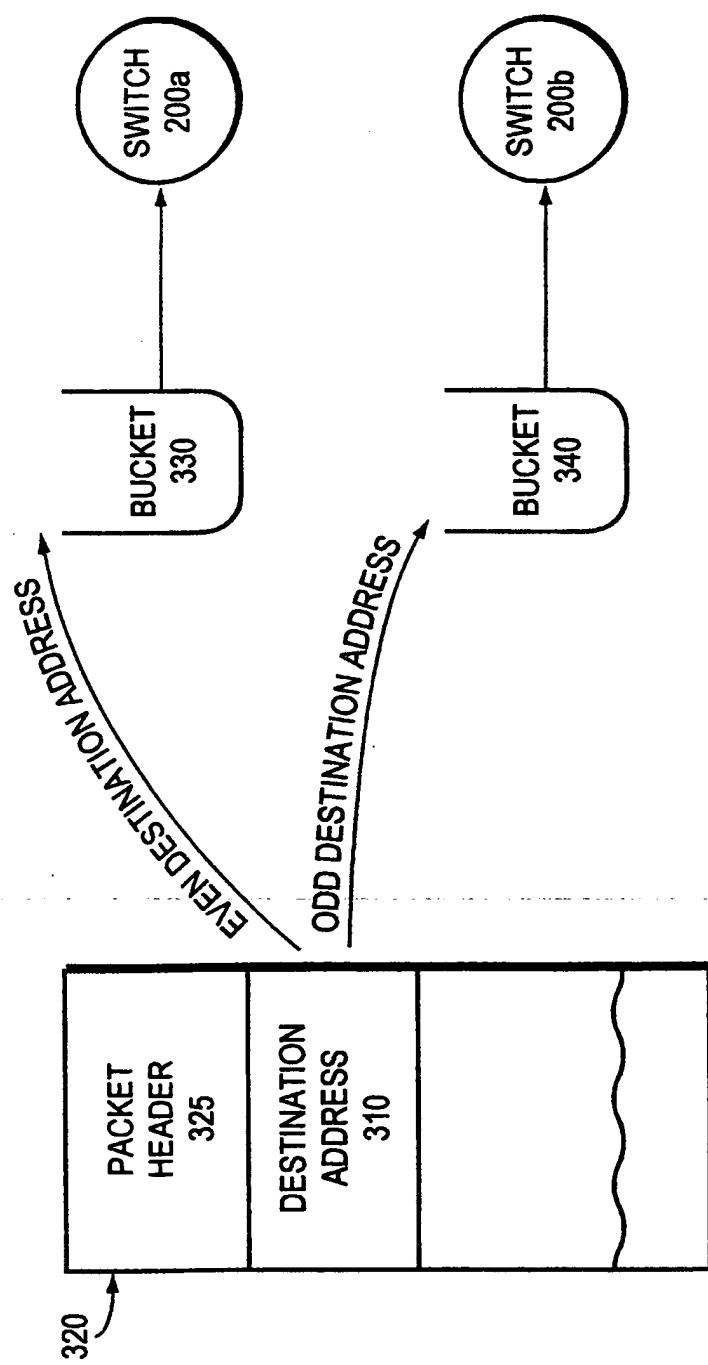
300

FIG. 3

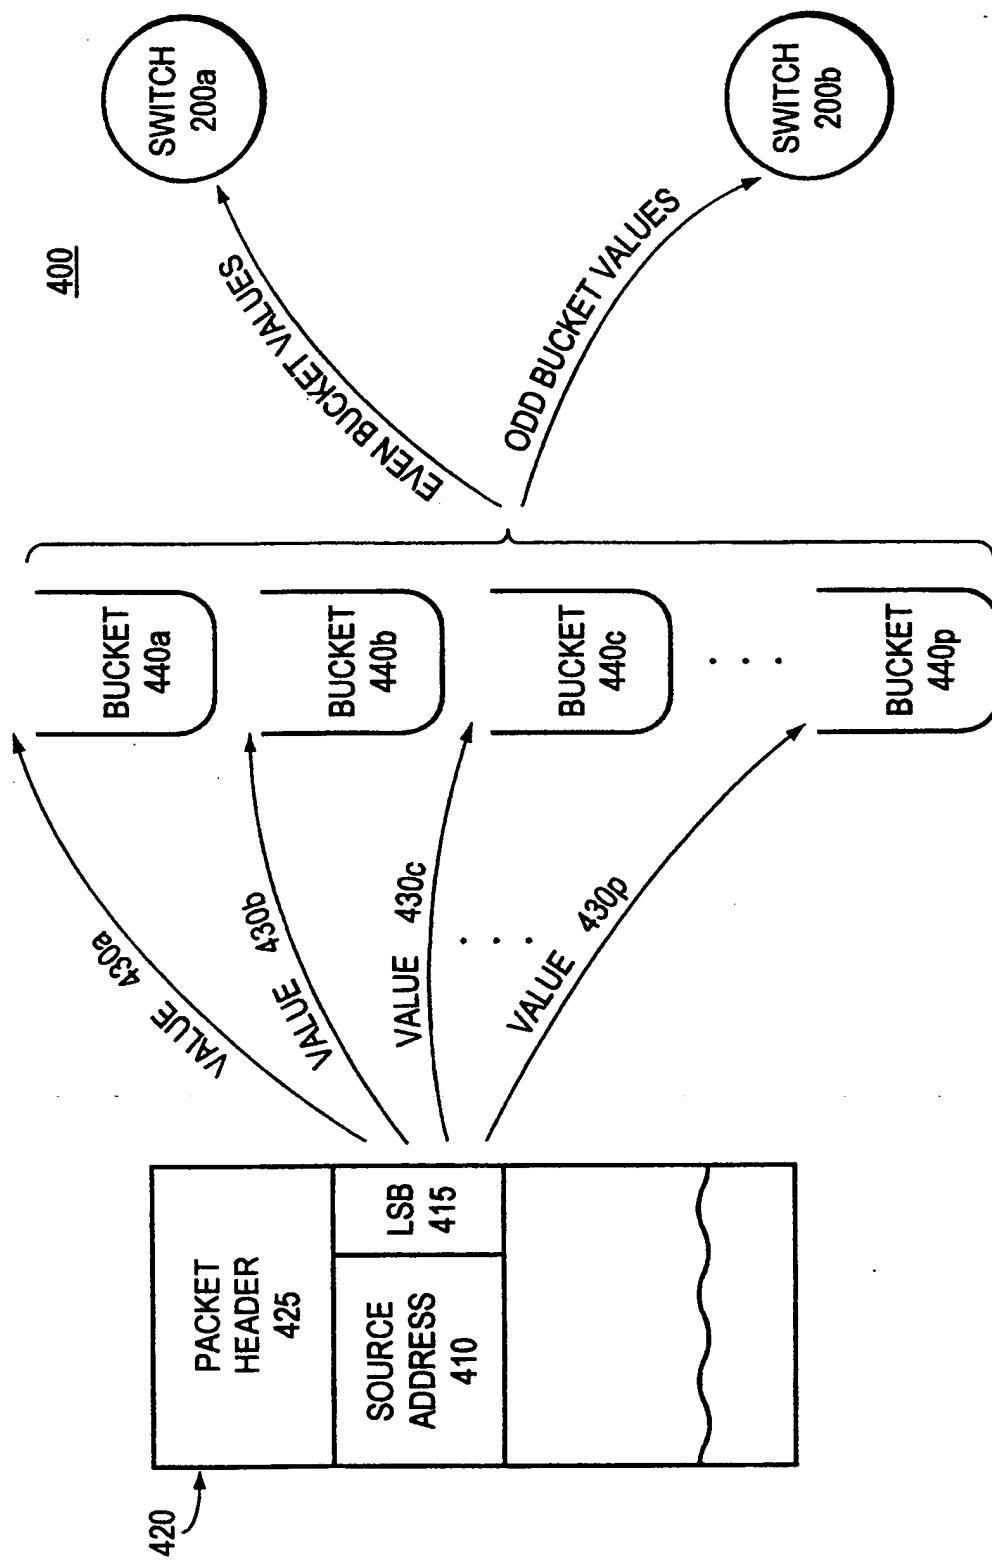


FIG. 4

5/6

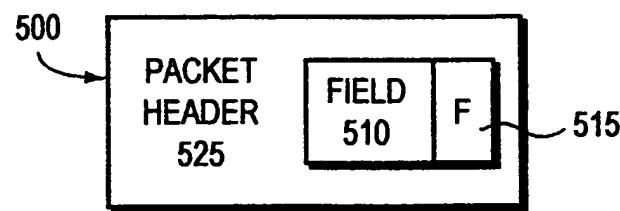


FIG. 5A

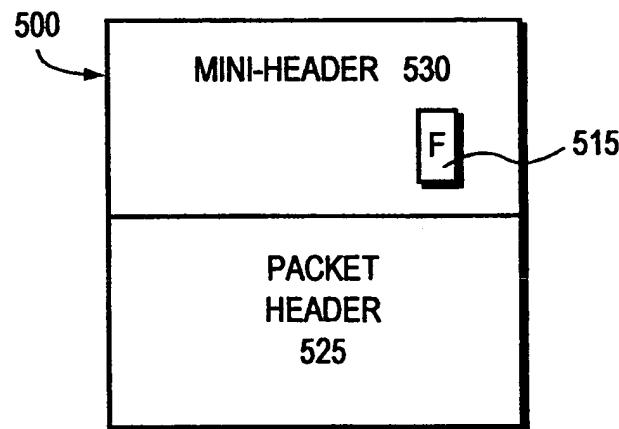


FIG. 5B

6/6

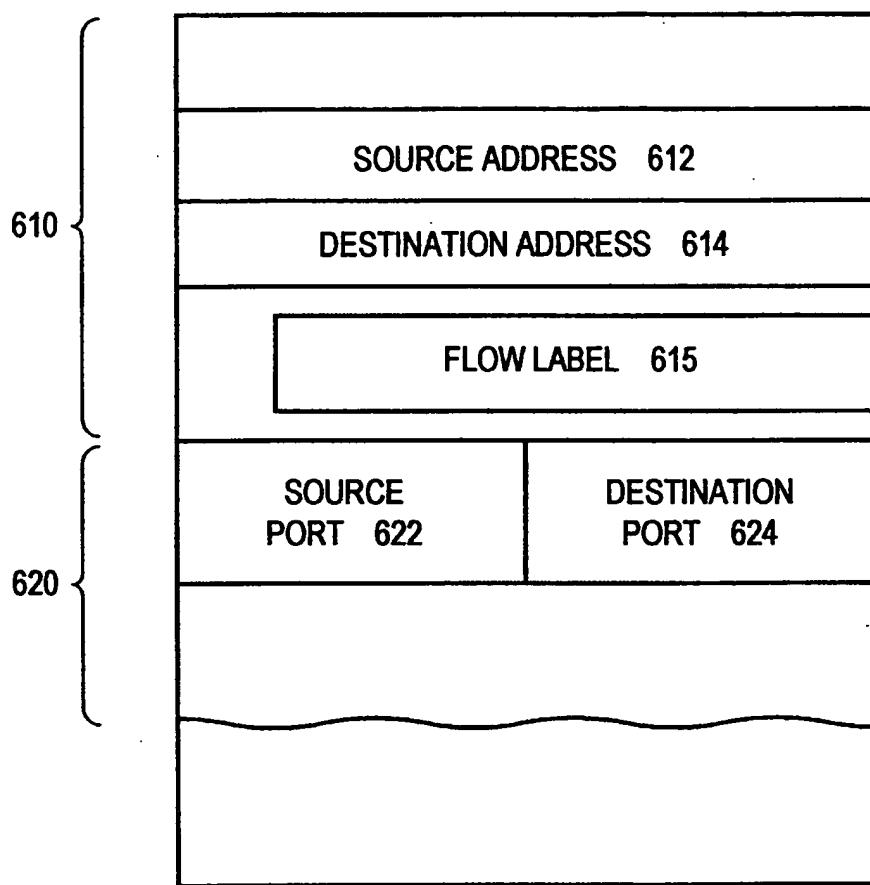
600

FIG. 6

INTERNATIONAL SEARCH REPORT

International Application No
PCT/US 99/10272

A. CLASSIFICATION OF SUBJECT MATTER
IPC 6 H04L29/06

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
IPC 6 H04L G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the International search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 4 621 359 A (MCMILLEN ROBERT J) 4 November 1986 (1986-11-04) figure 1 column 3, line 8 -column 4, line 41	1,17-19, 22,30
A		2-16,20, 21, 23-29, 31-34
A	EP 0 776 112 A (SUN MICROSYSTEMS INC) 28 May 1997 (1997-05-28) figures 1-7 column 5, line 23 -column 8, line 31	1,2, 20-23, 26,30,31

Further documents are listed in the continuation of box C.

Patent family members are listed in annex.

* Special categories of cited documents :

- "A" document defining the general state of the art which is not considered to be of particular relevance
- "E" earlier document but published on or after the International filing date
- "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- "O" document referring to an oral disclosure, use, exhibition or other means
- "P" document published prior to the International filing date but later than the priority date claimed

- "T" later document published after the International filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
- "&" document member of the same patent family

Date of the actual completion of the International search

28 September 1999

Date of mailing of the International search report

06/10/1999

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl
Fax: (+31-70) 340-3018

Authorized officer

Eraso Helguera, J

INTERNATIONAL SEARCH REPORT

Int'l Application No

PCT/US 99/10272

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 5 511 122 A (ATKINSON RANDALL) 23 April 1996 (1996-04-23) figure 4 column 10, line 19 -column 12, line 3	1,2,22, 23,26, 30,31
A	GB 2 301 919 A (BANKERS TRUST CO) 18 December 1996 (1996-12-18) figures 1,2 page 11, line 4 -page 15, line 4	1,2,22, 23,26, 30,31